
Clocks as Types

in Synchronous Dataflow Languages

Marc Pouzet
LRI & INRIA
Univ. Paris-Sud 11
Orsay

IFIP WG 2.8 – 9/06/2009

(joint work with Albert Cohen, Louis Mandel, Florence Plateau)

Synchronous Dataflow Languages

Model/program critical embedded software.

The idea of Lustre :

- ▶ directly write stream equations as **executable specifications**
- ▶ provide a **compiler** and associated analyzing tools to generate embedded code

E.g, the linear filter :

$$Y_0 = bX_0, \forall n Y_{n+1} = aY_n + bX_{n+1}$$

is programmed by writing, e.g :

$$Y = (0 \rightarrow a * \text{pre}(Y)) + Z;$$

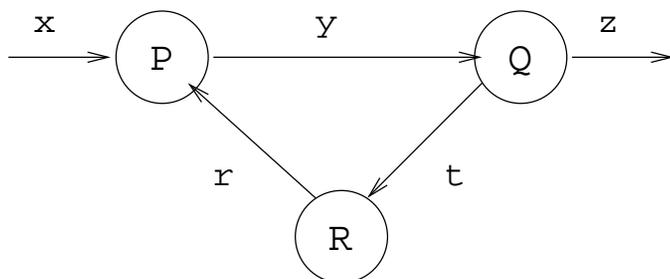
$$Z = b * X$$

we write **invariants**

- ▶ other primitives to deal with slow and fast processes (sub/over-sampling); not necessarily periodic

Dataflow Semantics

Kahn Principle : The semantics of process networks communicating through unbounded FIFOs (e.g., Unix pipe, sockets) ?



- message communication into FIFOs (send/wait)
- reliable channels, bounded communication delay
- blocking wait on a channel. The following program is **forbidden**

if (A is present) **or** (B is present) then ...

- a process = a continuous function $(V^\infty)^n \rightarrow (V'^\infty)^m$.

Lustre :

- Lustre has a **Kahn semantics** (no test of absence)
- A dedicated **type system** (clock calculus) to guaranty the existence of an execution with no buffer (no synchronization)

Pros and Cons of KPN

(+) : **Simple semantics** : a process defines a function (determinism);
composition is function composition

(+) : **Modularity** : a network is a continuous function

(+) : **Asynchronous distributed execution** : easy; no centralized scheduler

(+/-) : **Time invariance** : no explicit timing; but impossible to state that two events happen at the same time.

x	=	x_0		x_1		x_2		x_3	x_4	x_5		...
$f(x)$	=	y_0		y_1		y_2		y_3	y_4	y_5		...
$f(x)$	=	y_0		y_1	y_2			y_3		y_4	y_5	...

This appeared to be a useful model for video apps (TV boxes) : Sally (Philips NatLabs), StreamIt (MIT), Xstream (ST-micro) with various “synchronous” restriction *à la SDF* (Edward Lee)

A small dataflow kernel

A small kernel with minimal primitives

$$\begin{aligned} e \quad ::= \quad & e \text{ fby } e \mid op(e, \dots, e) \mid x \mid i \\ & \mid \text{merge } e \ e \ e \mid e \text{ when } e \\ & \mid \lambda x.e \mid e \ e \mid \text{rec } x.e \\ op \quad ::= \quad & + \mid - \mid \text{not} \mid \dots \end{aligned}$$

- function ($\lambda x.e$), application ($e \ e$), fix-point ($\text{rec } x.e$)
- constants i and variables (x)
- dataflow primitives : $x \text{ fby } y$ is the unitary delay ; $op(e_1, \dots, e_n)$ the point-wise application ; sub-sampling/oversampling (when/merge).

Dataflow Primitives

x	x_0	x_1	x_2	x_3	x_4	x_5
y	y_0	y_1	y_2	y_3	y_4	y_5
$x + y$	$x_0 + y_0$	$x_1 + y_1$	$x_2 + y_2$	$x_3 + y_3$	$x_4 + y_4$	$x_5 + y_5$
$x \text{ fby } y$	x_0	y_0	y_1	y_2	y_3	y_4
h	1	0	1	0	1	0
$x' = x \text{ when } h$	x_0		x_2		x_4	
z		z_0		z_1		z_2
$\text{merge } h \ x' \ z$	x_0	z_0	x_2	z_1	x_4	z_2

Sampling :

- ▶ if h is a boolean sequence, $x \text{ when } h$ produces a sub-sequence of x
- ▶ $\text{merge } h \ x \ z$ combines two sub-sequences

Kahn Semantics

Every operator is interpreted as a stream function ($V^\infty = V^* + V^\omega$). E.g., if $x \mapsto s_1$ and $y \mapsto s_2$ then the value of $x + y$ is $+^\# (s_1, s_2)$

$$i^\# = i.i^\#$$

$$+^\# (x.s_1, y.s_2) = (x + y).+^\# (s_1, s_2)$$

$$(x.s_1) \text{ fby}^\# s_2 = x.s_2$$

$$x.s \text{ when}^\# 1.c = x.(s \text{ when}^\# c)$$

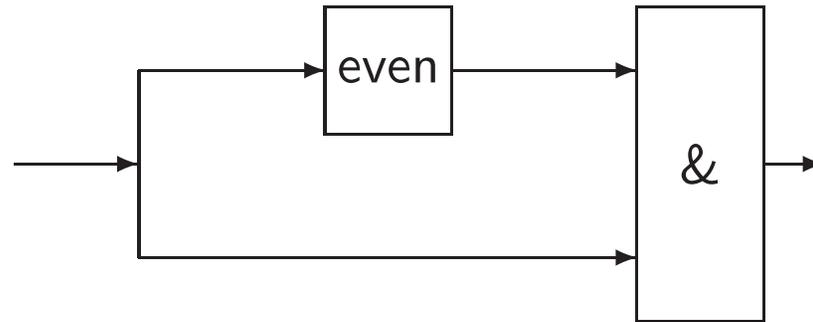
$$x.s \text{ when}^\# 0.c = s \text{ when}^\# c$$

$$\text{merge}^\# 1.c x.s_1 s_2 = x.\text{merge}^\# c s_1 s_2$$

$$\text{merge}^\# 0.c s_1 y.s_2 = y.\text{merge}^\# c s_1 s_2$$

Synchrony

Some programs generate monsters.



If $x = (x_i)_{i \in \mathbb{N}}$ then $\text{even}(x) = (x_{2i})_{i \in \mathbb{N}}$ and $x \& \text{even}(x) = (x_i \& x_{2i})_{i \in \mathbb{N}}$.

Unbounded FIFOs !

- ▶ must be rejected statically
- ▶ every operator is finite memory through the composition is not : all the complexity (synchronization) is hidden in communication channels
- ▶ the Kahn semantics does not model time, i.e., impossible to state that two event arrive **at the same time**

Synchronous (Clocked) streams

Complete streams with an explicit representation of absence (*abs*).

$$x : (V^{abs})^\infty$$

Clock : the clock of x is a boolean sequence

$$B = \{0, 1\}$$

$$CLOCK = B^\infty$$

$$\text{clock } \epsilon = \epsilon$$

$$\text{clock } (abs.x) = 0.\text{clock } x$$

$$\text{clock } (v.x) = 1.\text{clock } x$$

Synchronous streams :

$$ClStream(V, cl) = \{s/s \in (V^{abs})^\infty \wedge \text{clock } s \leq_{prefix} cl\}$$

An other possible encoding : $x : (V \times \mathbb{N})^\infty$

Dataflow Primitives

Constant :

$$i^\#(\epsilon) = \epsilon$$

$$i^\#(1.cl) = i.i^\#(cl)$$

$$i^\#(0.cl) = abs.i^\#(cl)$$

Point-wise application :

Synchronous arguments must be constant, i.e., having the same clock

$$+^\#(s_1, s_2) = \epsilon \text{ if } s_i = \epsilon$$

$$+^\#(abs.s_1, abs.s_2) = abs.+^\#(s_1, s_2)$$

$$+^\#(v_1.s_1, v_2.s_2) = (v_1 + v_2).+^\#(s_1, s_2)$$

Partial definitions

What happens when one element is present and the other is absent ?

Constraint their domain :

$(+) : \forall cl : \mathcal{CLOCK}. ClStream(\mathbf{int}, cl) \times ClStream(\mathbf{int}, cl) \rightarrow ClStream(\mathbf{int}, cl)$

i.e., $(+)$ expect its two input stream to be on the same clock cl and produce an output on the same clock

These extra conditions are **types** which must be statically verified

Remark (notation) : Regular types and clock types can be written separately :

- $(+) : \mathbf{int} \times \mathbf{int} \rightarrow \mathbf{int}$ ← **its type**
- $(+) :: \forall cl. cl \times cl \rightarrow cl$ ← **its clock type**

In the following, we only consider the clock type.

Sampling

$$s_1 \text{ when}^\# s_2 = \epsilon \text{ if } s_1 = \epsilon \text{ or } s_2 = \epsilon$$

$$(abs.s) \text{ when}^\# (abs.c) = abs.s \text{ when}^\# c$$

$$(v.s) \text{ when}^\# (1.c) = v.s \text{ when}^\# c$$

$$(v.s) \text{ when}^\# (0.c) = abs.x \text{ when}^\# c$$

$$\text{merge } c \ s_1 \ s_2 = \epsilon \text{ if one of the } s_i = \epsilon$$

$$\text{merge } (abs.c) \ (abs.s_1) \ (abs.s_2) = abs.\text{merge } c \ s_1 \ s_2$$

$$\text{merge } (1.c) \ (v.s_1) \ (abs.s_2) = v.\text{merge } c \ s_1 \ s_2$$

$$\text{merge } (0.c) \ (abs.s_1) \ (v.s_2) = v.\text{merge } c \ s_1 \ s_2$$

Examples

$base = (1)$	1	1	1	1	1	1	1	1	1	1	1	1	...
x	x_0	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}	...
$h = (10)$	1	0	1	0	1	0	1	0	1	0	1	0	...
$y = x \text{ when } h$	x_0		x_2		x_4		x_6		x_8		x_{10}	x_{11}	...
$h' = (100)$	1		0		0		1		0		0	1	...
$z = y \text{ when } h'$	x_0						x_6					x_{11}	...
k			k_0		k_1				k_2		k_3		...
$\text{merge } h' \ z \ k$	x_0		k_0		k_1		x_6		k_2		k_3		...

let clock five =

let rec f = true fby false fby false fby false fby f in f

let node stutter x = o where

rec o = merge five x ((0 fby o) whenot five) in o

$\text{stutter}(\text{nat}) = 0.0.0.0.1.1.1.1.2.2.2.2.3.3\dots$

Sampling and clocks

- ▶ $x \text{ when}^\# y$ is defined when x and y have the same clock cl
- ▶ the clock of $x \text{ when}^\# c$ is written $cl \text{ on } c$: “ c moves at the pace of cl ”

$$\begin{aligned} s \text{ on } c &= \epsilon \text{ if } s = \epsilon \text{ or } c = \epsilon \\ (1.cl) \text{ on } (1.c) &= 1.cl \text{ on } c \\ (1.cl) \text{ on } (0.c) &= 0.cl \text{ on } c \\ (0.cl) \text{ on } (abs.c) &= 0.cl \text{ on } c \end{aligned}$$

We get :

$$\text{when} : \forall cl. \forall x : cl. \forall c : cl. cl \text{ on } c$$

$$\text{merge} : \forall cl. \forall c : cl. \forall x : cl \text{ on } c. \forall y : cl \text{ on } not\ c. cl$$

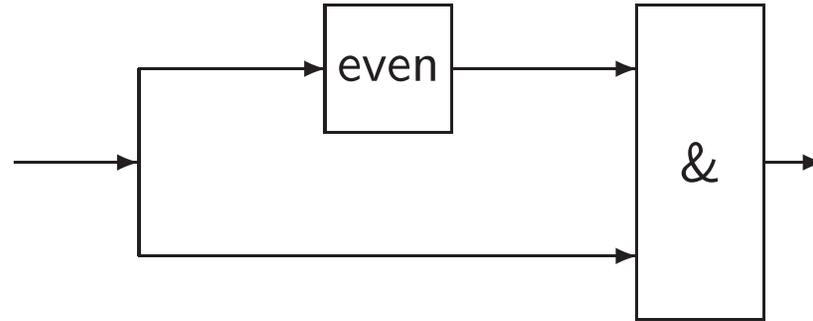
Written instead :

$$\text{when} : \forall cl. cl \rightarrow (c : cl) \rightarrow cl \text{ on } c$$

$$\text{merge} : \forall cl. (c : cl) \rightarrow cl \text{ on } c \rightarrow cl \text{ on } not\ c \rightarrow cl$$

Checking Synchrony

The previous program is now rejected.



This is now a **typing error**

```
let even x = x when half
let non_synchronous x = x & (even x)
                        ^^^^^^^
```

This expression has clock 'a on half,
but is used with clock 'a

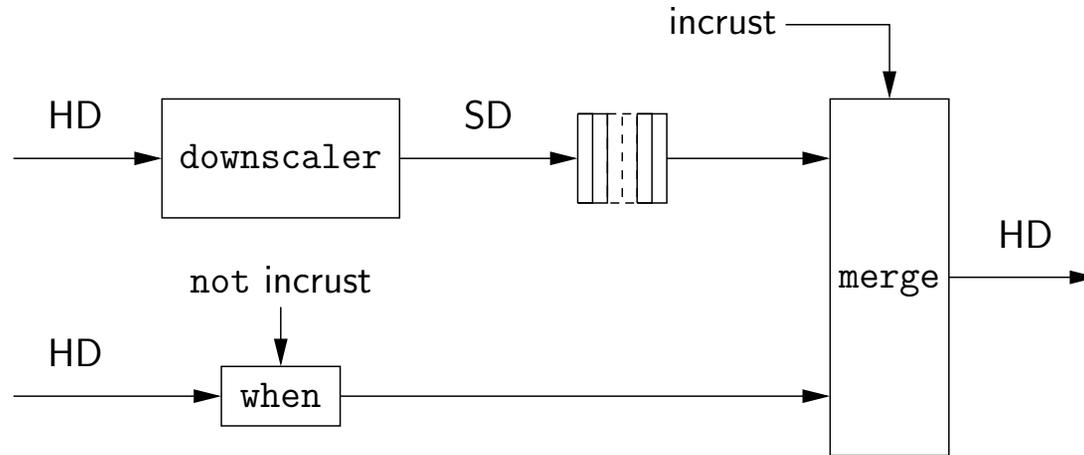
Final remarks :

- We only considered **clock equality**, i.e., “two streams are either synchronous or not”
- Clocks are used extensively to generate **efficient sequential code**

From Synchrony to Relaxed Synchrony

- can we compose non strictly synchronous streams provided their clocks are closed from each other ?
- communication between systems which are “almost” synchronous
- model jittering, bounded delays
- Give more freedom to the compiler, generate more efficient code, translate into regular synchronous code if necessary

A typical example : Picture in Picture



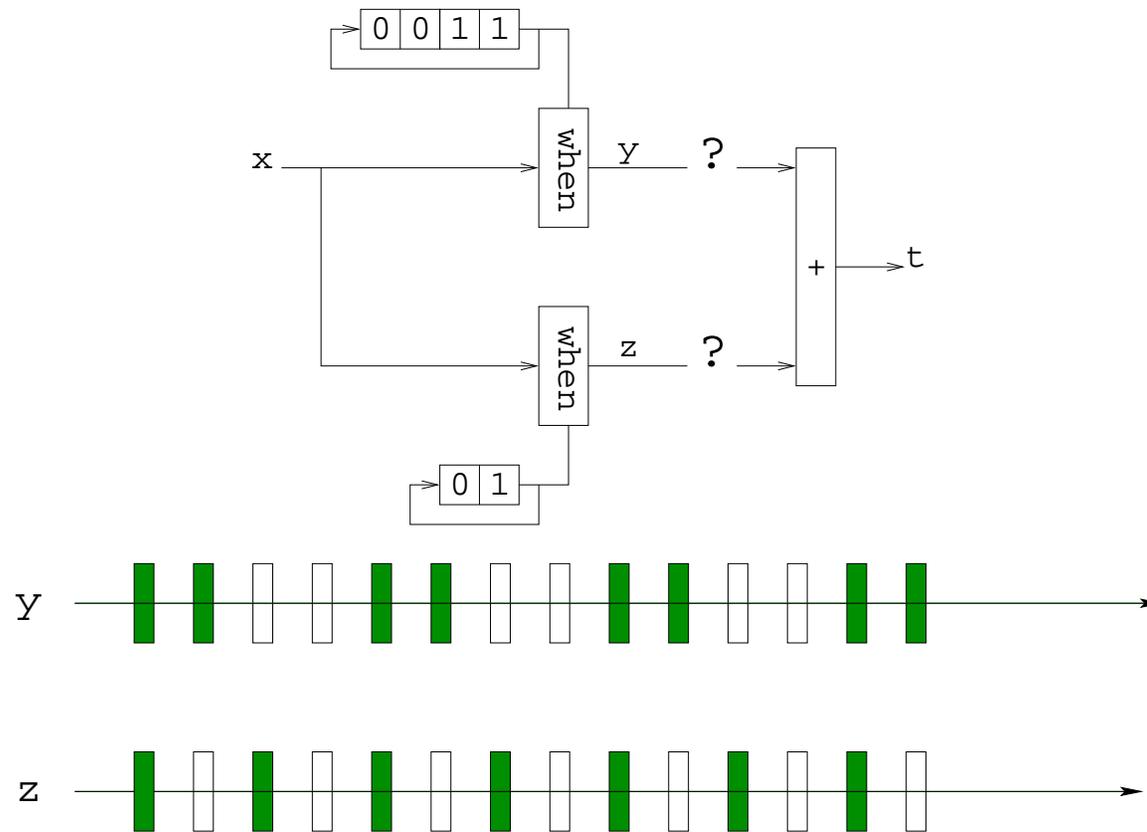
Incrustation of a Standard Definition (SD) image in a High Definition (HD) one

- ▶ **downscaler** : reduction of an HD image (1920×1080 pixels) to an SD image (720×480 pixels)
- ▶ **when** : removal of a part of an HD image
- ▶ **merge** : incrustation of an SD image in an HD image

Question :

- ▶ buffer size needed between the downscaler and the merge nodes ?
- ▶ delay introduced by the picture in picture in the video processing chain ?

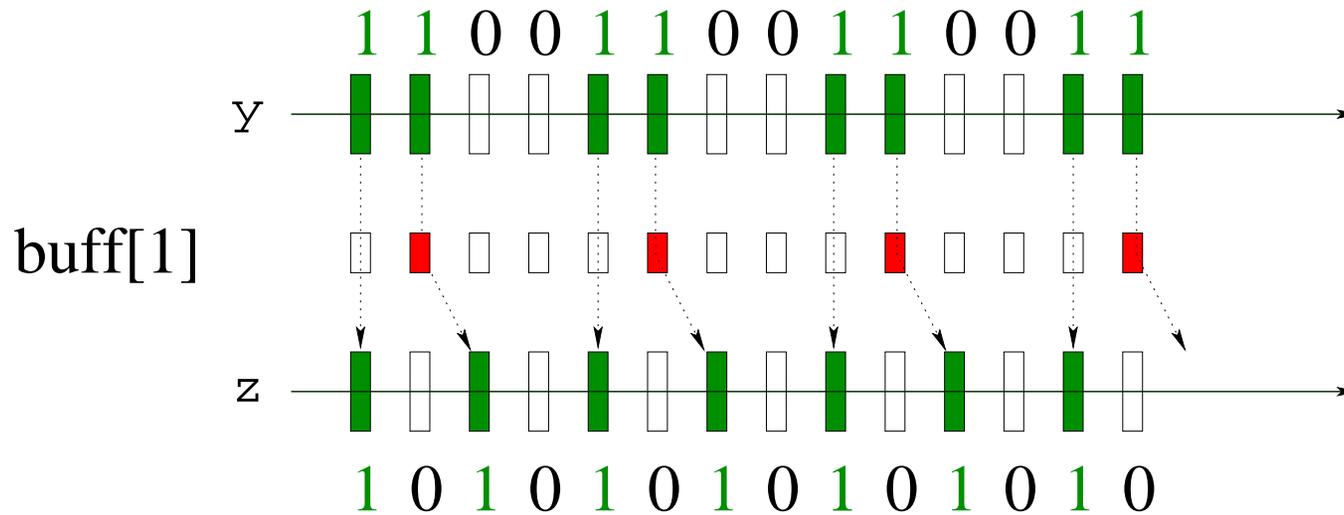
Too restrictive for video applications



- ▶ streams should be synchronous
- ▶ adding buffer (by hand) difficult and error-prone
- ▶ compute it automatically and generate synchronous code

relax the associated clocking rules

N -Synchronous Kahn Networks



- based on the use of *infinite ultimately periodic sequences*
- a precedence relation $cl_1 <: cl_2$

Ultimately periodic sequences

\mathbb{Q}_2 for the set of infinite periodic binary words.

$$(01) = 01\ 01\ 01\ 01\ 01\ 01\ 01\ 01\ 01\ \dots$$

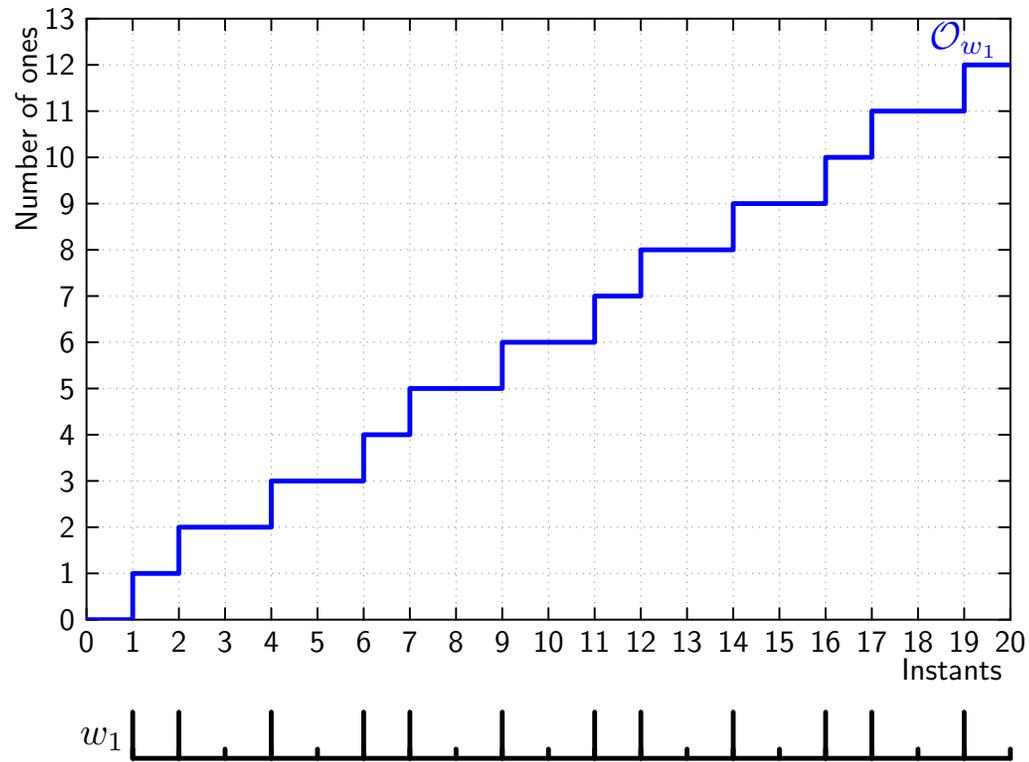
$$0(1101) = 0\ 1101\ 1101\ 1101\ 1101\ 1101\ 1101\ 1101\ \dots$$

- 1 for presence
- 0 for absence

Definition :

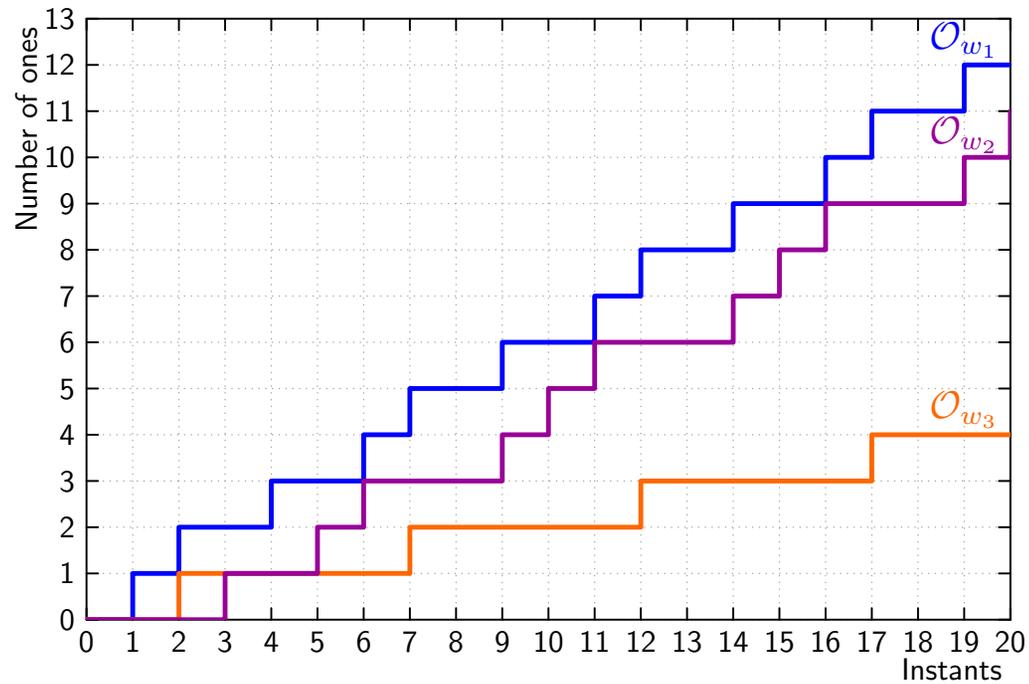
$$w ::= u(v) \quad \text{where } u \in (0 + 1)^* \text{ and } v \in (0 + 1)^+$$

Clocks and infinite binary words



$\mathcal{O}_w(i) = \text{cumulative function of 1 from } w$

Clocks and infinite binary words



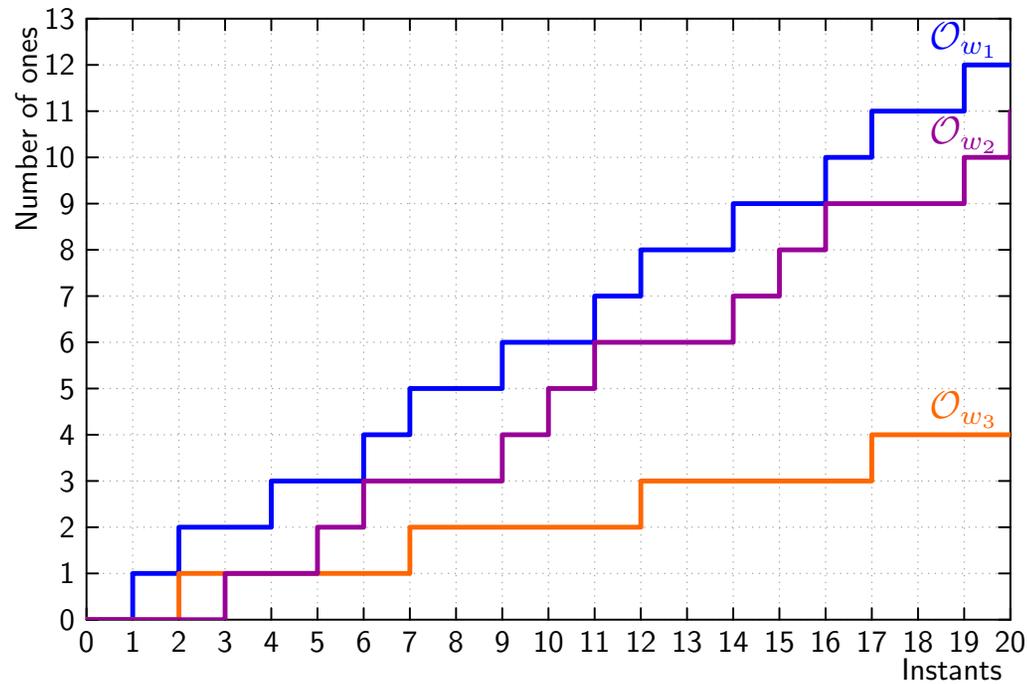
buffer

$$size(w_1, w_2) = \max_{i \in \mathbb{N}} (\mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i))$$

sub-typing

$$w_1 <: w_2 \stackrel{def}{\iff} \exists n \in \mathbb{N}, \forall i, 0 \leq \mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i) \leq n$$

Clocks and infinite binary words



buffer

$$size(w_1, w_2) = \max_{i \in \mathbb{N}} (\mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i))$$

sub-typing

$$w_1 <: w_2 \stackrel{def}{\iff} \exists n \in \mathbb{N}, \forall i, 0 \leq \mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i) \leq n$$

synchronizability

$$w_1 \bowtie w_2 \stackrel{def}{\iff} \exists b_1, b_2 \in \mathbb{Z}, \forall i, b_1 \leq \mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i) \leq b_2$$

precedence

$$w_1 \preceq w_2 \stackrel{def}{\iff} \forall i, \mathcal{O}_{w_1}(i) \geq \mathcal{O}_{w_2}(i)$$

Multi-clock

$$c ::= w \mid c \text{ on } w \quad w \in (0 + 1)^\omega$$

c on w is a **sub-clock** of c , by moving in w at the pace of c . E.g.,
 $1(10)$ on $(01) = (0100)$.

base	1 1 1 1 1 1 1 1 1 1 1 ...	(1)
p_1	1 1 0 1 0 1 0 1 0 1 ...	1(10)
base on p_1	1 1 0 1 0 1 0 1 0 1 ...	1(10)
p_2	0 1 0 1 0 1 ...	(01)
(base on p_1) on p_2	0 1 0 0 0 1 0 0 0 1 ...	(0100)

For ultimately periodic clocks, precedence, synchronizability and equality are decidable (but expensive)

Come-back to the language

Pure synchrony :

- ▶ close to an ML type system (e.g., SCADE 6)
- ▶ structural equality of clocks

$$\frac{H \vdash e_1 : ck \quad H \vdash e_2 : ck}{H \vdash op(e_1, e_2) : ck}$$

Relaxed Synchrony :

- ▶ we add a **sub-typing** rule :

$$\text{(SUB) } \frac{H \vdash e : ck \text{ on } w \quad w <: w'}{H \vdash e : ck \text{ on } w'}$$

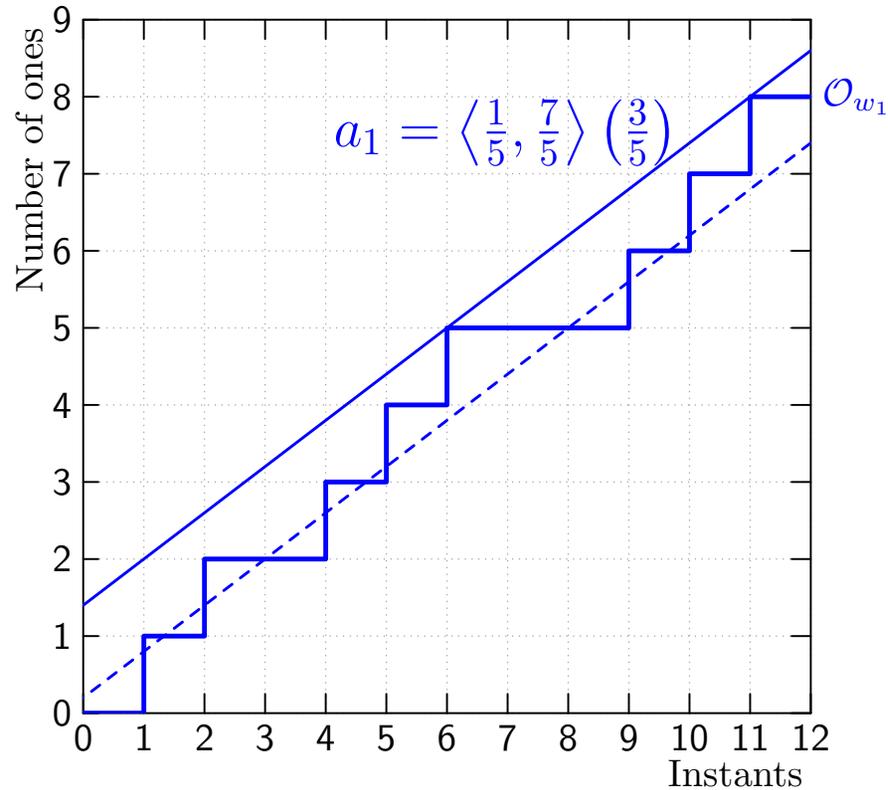
- ▶ defines synchronization points when a buffer is inserted

What about non periodic systems ?

- ▶ The same idea : synchrony + properties between clocks. Insuring the absence of deadlocks and bounded buffering.
- ▶ The **exact** computation with periodic clocks does not work in practice (and is useless). E.g., (10100100) on $0^{3600}(1)$ on $(101001001) = 0^{9600}(10^4 10^7 10^7 10^2)$
- ▶ Motivations :
 1. To treat long periodic patterns. To avoid an exact computation.
 2. To deal with almost periodic clocks. E.g., α on w where $w = 00.((10) + (01))^*$
(e.g. $w = 00 01 10 01 01 10 01 10 \dots$)

Idea : manipulate sets of clocks ; turn questions into arithmetic ones

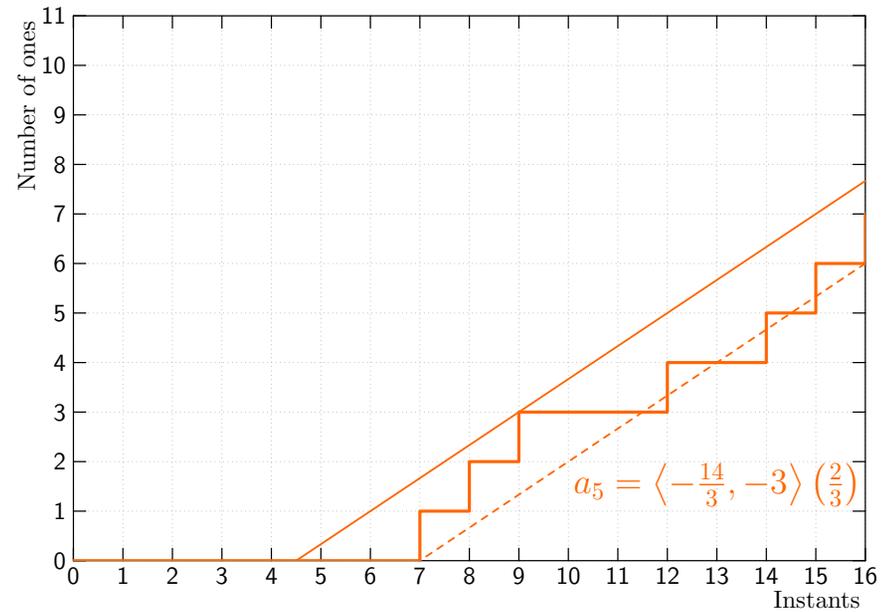
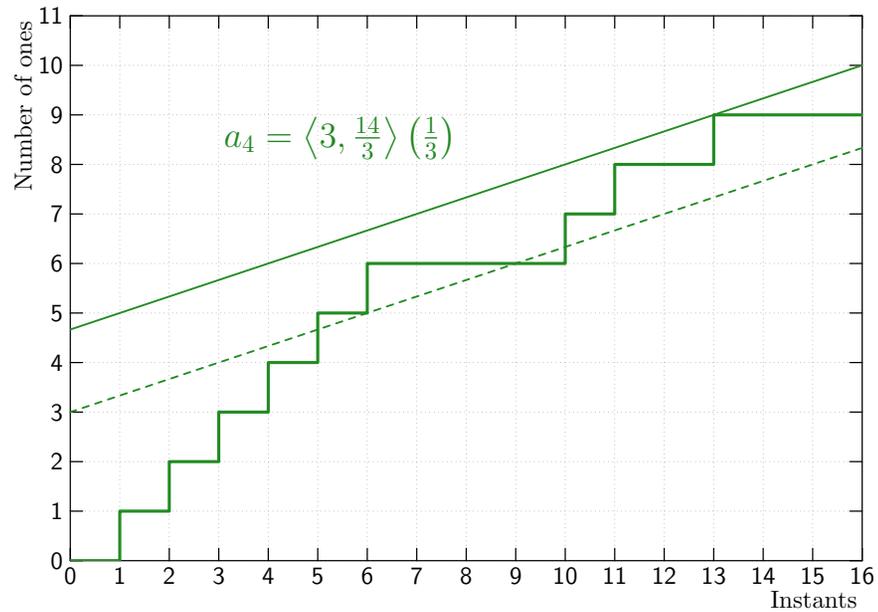
Abstraction of Infinite Binary Words



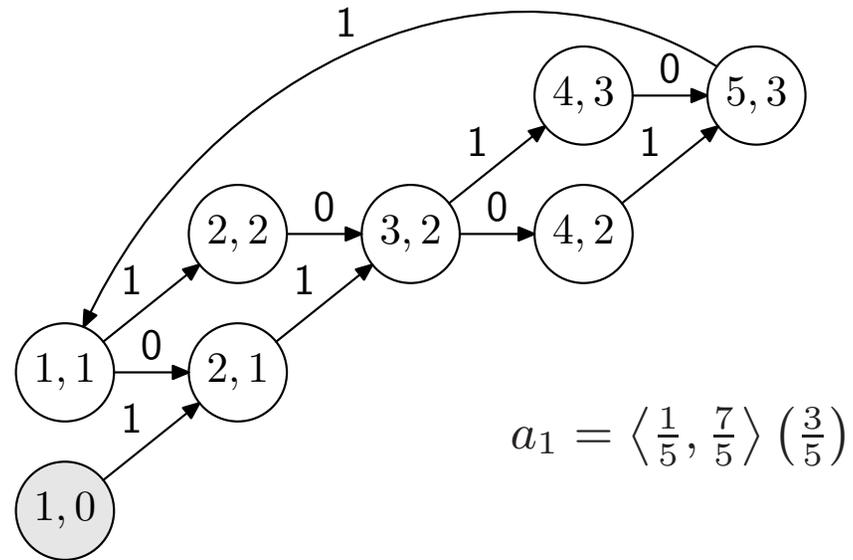
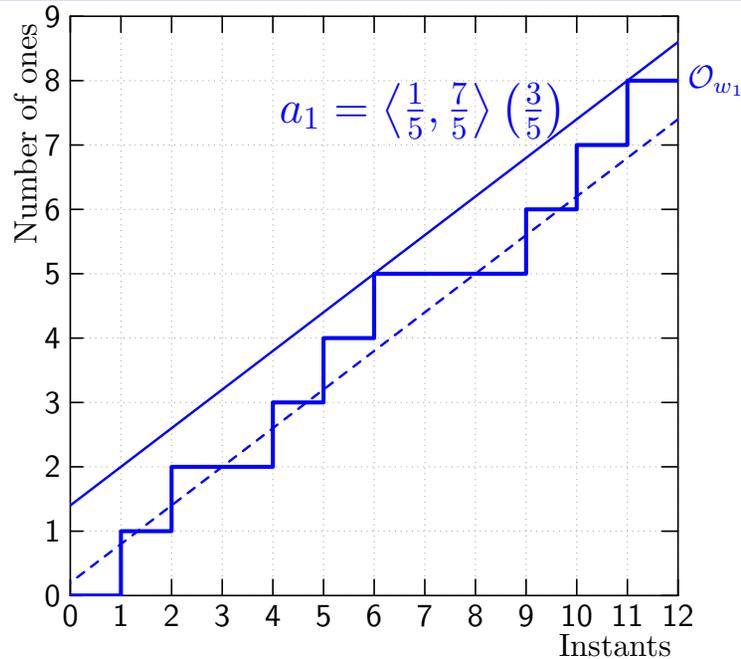
A word w can be abstracted by two lines : $abs(w) = \langle b^0, b^1 \rangle (r)$

$$concr \left(\langle b^0, b^1 \rangle (r) \right) \stackrel{def}{\Leftrightarrow} \left\{ w, \forall i \geq 1, \wedge \begin{array}{l} w[i] = 1 \Rightarrow O_w(i) \leq r \times i + b^1 \\ w[i] = 0 \Rightarrow O_w(i) \geq r \times i + b^0 \end{array} \right\}$$

Abstraction of Infinite Binary Words



Abstract Clocks as Automata



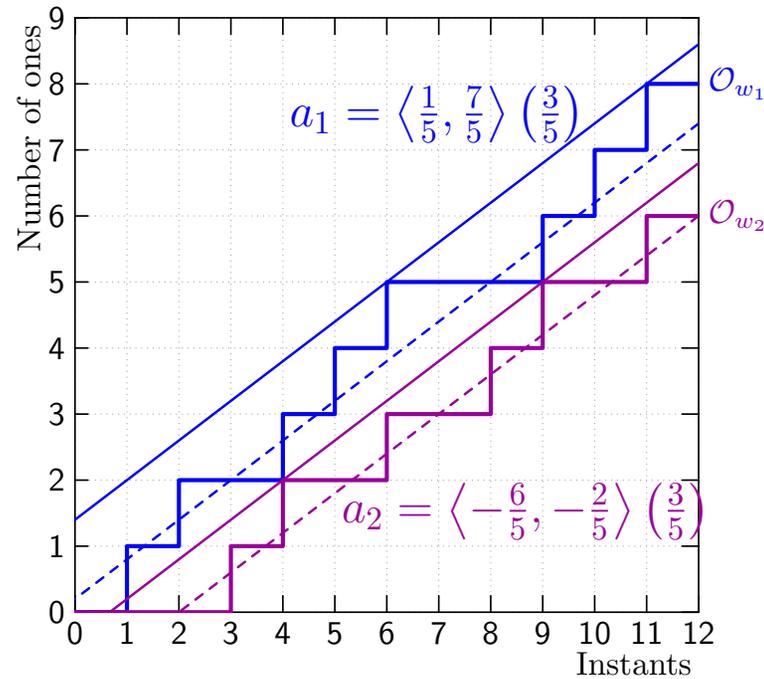
► set of states $\{(i, j) \in \mathbb{N}^2\}$: coordinates in the 2D-chronogram

► finite number of state equivalence classes

► transition function δ :
$$\begin{cases} \delta(1, (i, j)) = nf(i + 1, j + 1) & \text{if } j + 1 \leq r \times i + b^1 \\ \delta(0, (i, j)) = nf(i + 1, j + 0) & \text{if } j + 0 \geq r \times i + b^0 \end{cases}$$

► allows to check/generate clocks

Abstract Relations



Synchronizability : $r_1 = r_2 \Leftrightarrow \langle b^0_1, b^1_1 \rangle (r_1) \bowtie^{\sim} \langle b^0_2, b^1_2 \rangle (r_2)$

Precedence : $b^1_2 - b^0_1 < 1 \Rightarrow \langle b^0_1, b^1_1 \rangle (r) \preceq^{\sim} \langle b^0_2, b^1_2 \rangle (r)$

Subtyping : $a_1 <:\sim a_2 \Leftrightarrow a_1 \bowtie^{\sim} a_2 \wedge a_1 \preceq^{\sim} a_2$

▷ proposition : $abs(w_1) <:\sim abs(w_2) \Rightarrow w_1 <:\sim w_2$

▷ buffer : $size(a_1, a_2) = \lfloor b^1_1 - b^0_2 \rfloor$

Abstract Operators

Composed clocks : $c ::= w \mid \mathit{not} w \mid c \mathit{on} c$

Abstraction of a composed clock :

$$\mathit{abs}(\mathit{not} w) = \mathit{not}^{\sim} \mathit{abs}(w)$$

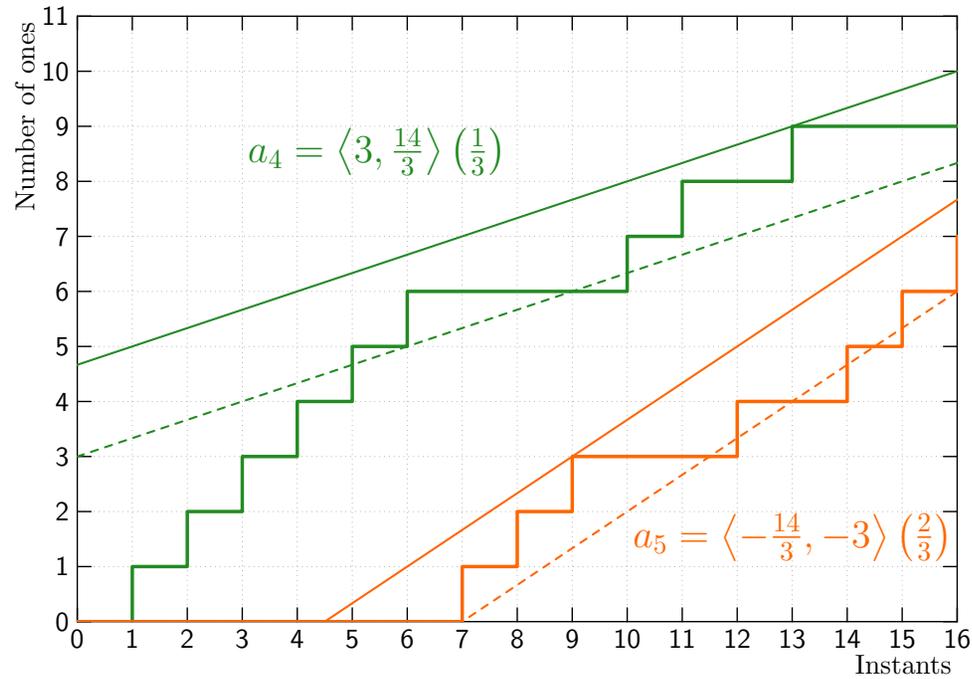
$$\mathit{abs}(c_1 \mathit{on} c_2) = \mathit{abs}(c_1) \mathit{on}^{\sim} \mathit{abs}(c_2)$$

Operators correctness property :

$$\mathit{not} w \in \mathit{concr}(\mathit{not}^{\sim} \mathit{abs}(w))$$

$$c_1 \mathit{on} c_2 \in \mathit{concr}(\mathit{abs}(c_1) \mathit{on}^{\sim} \mathit{abs}(c_2))$$

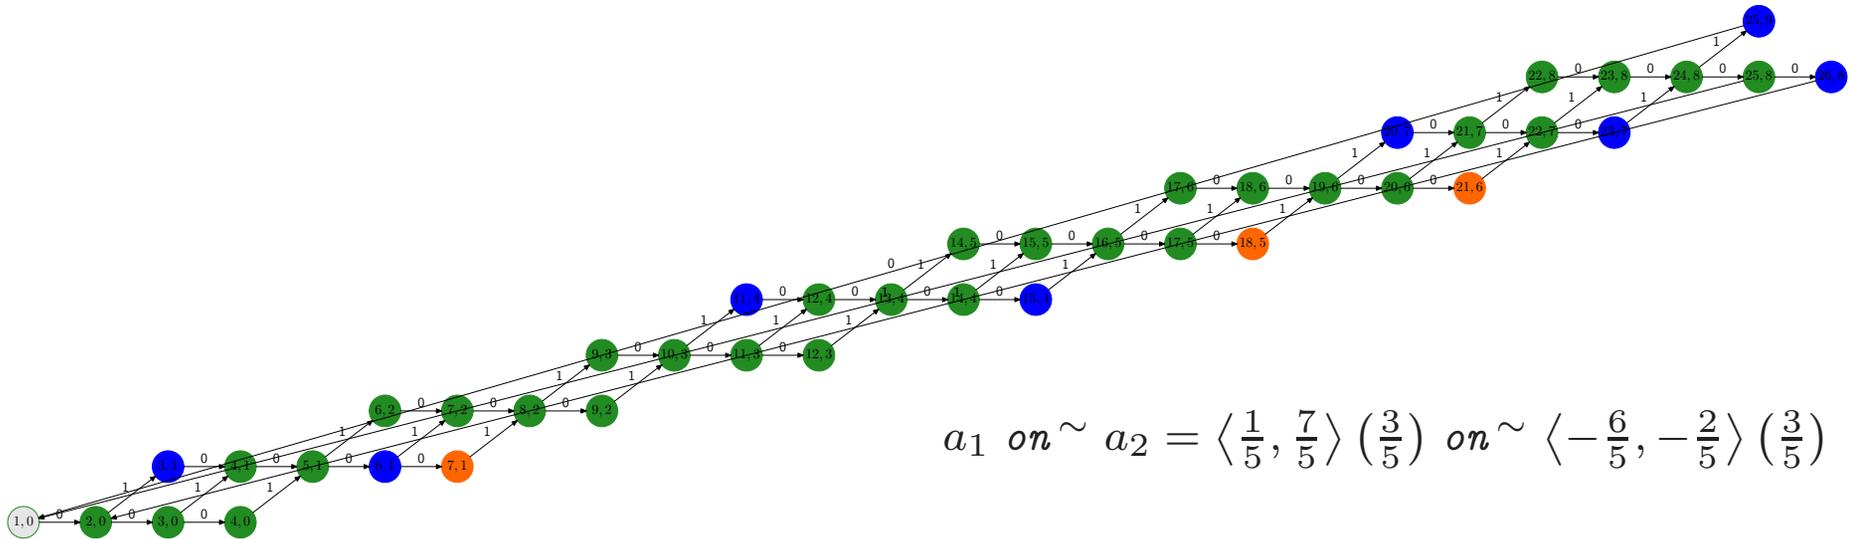
Abstract Operators



not^{\sim} operator definition :

$$\blacktriangleright not^{\sim} \langle b^0, b^1 \rangle (r) = \langle -b^1, -b^0 \rangle (1 - r)$$

Abstract Operators



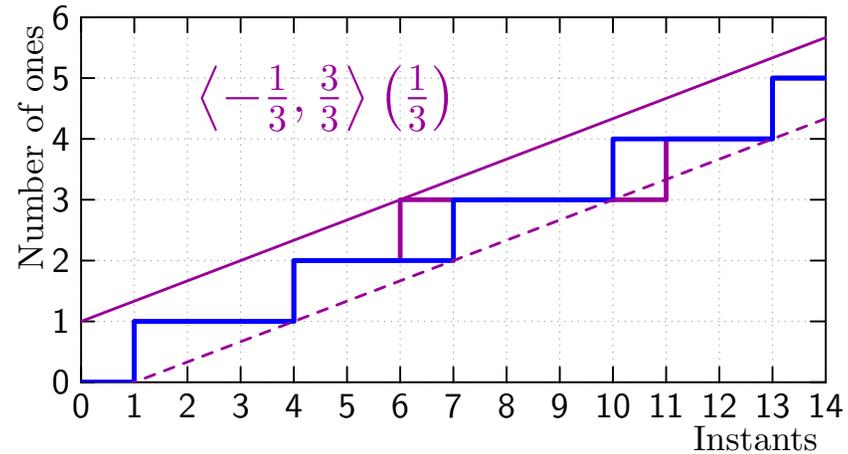
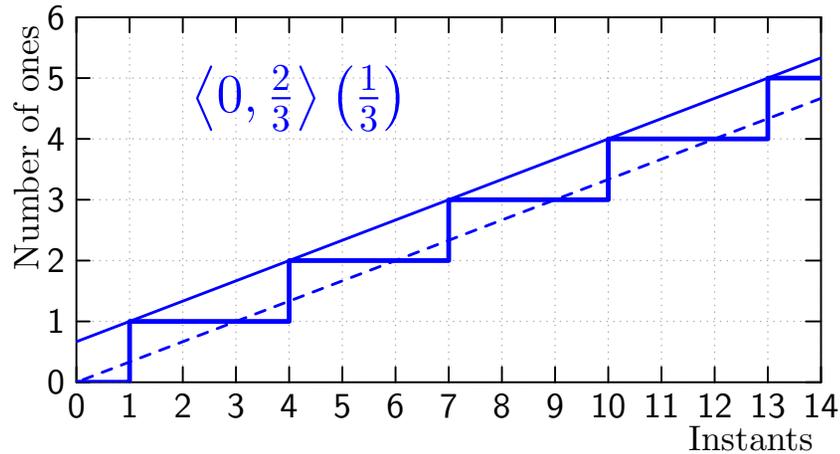
$$a_1 \text{ on}^{\sim} a_2 = \left\langle \frac{1}{5}, \frac{7}{5} \right\rangle \left(\frac{3}{5} \right) \text{ on}^{\sim} \left\langle -\frac{6}{5}, -\frac{2}{5} \right\rangle \left(\frac{3}{5} \right)$$

on^{\sim} operator definition :

$$\begin{aligned} & \left\langle b^0_1, b^1_1 \right\rangle \left(r_1 \right) \\ \text{on}^{\sim} & \left\langle b^0_2, b^1_2 \right\rangle \left(r_2 \right) \\ = & \left\langle b^0_1 \times r_2 + b^0_2, b^1_1 \times r_2 + b^1_2 \right\rangle \left(r_1 \times r_2 \right) \end{aligned}$$

with $b^0_1 \leq 0, b^0_2 \leq 0$

Modeling Jitter



- ▶ set of clock of rate $r = \frac{1}{3}$ and jitter 1 can be specified by $\langle -\frac{1}{3}, \frac{3}{3} \rangle (\frac{1}{3})$
- ▶ $\langle -\frac{1}{3}, \frac{3}{3} \rangle (\frac{1}{3}) = \langle -1, 1 \rangle (1) \text{ on } \sim \langle 0, \frac{2}{3} \rangle (\frac{1}{3})$
- ▶ $f :: \forall \alpha. \alpha \rightarrow \alpha \text{ on } \sim \langle -\frac{1}{3}, \frac{3}{3} \rangle (\frac{1}{3})$

Formalization in a Proof Assistant

Most of the properties have been proved in Coq

- ▶ example of property

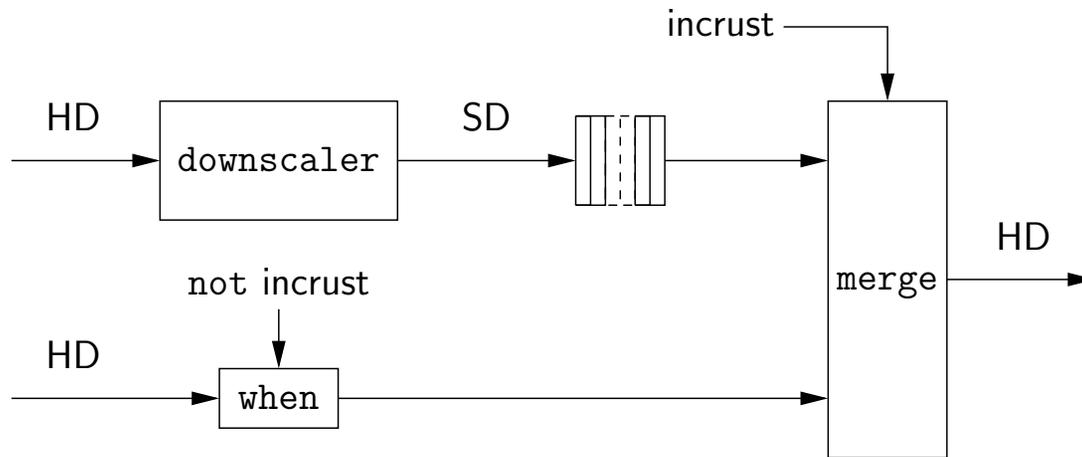
```
Property on_absh_correctness:
```

```
  forall (w1:ibw) (w2:ibw),  
  forall (a1:abstractionh) (a2:abstractionh),  
  forall H_wf_a1: well_formed_abstractionh a1,  
  forall H_wf_a2: well_formed_abstractionh a2,  
  forall H_a1_eq_absh_w1: in_abstractionh w1 a1,  
  forall H_a2_eq_absh_w2: in_abstractionh w2 a2,  
  in_abstractionh (on w1 w2) (on_absh a1 a2).
```

- ▶ number of Source Lines of Code

- ▶ specifications : about 1600 SLOC
- ▶ proofs : about 5000 SLOC

Back to the Picture in Picture Example



- ▶ abstraction of downscaler output :

$$\begin{aligned}
 & \text{abs}((10100100) \text{ on } 0^{3600}(1) \text{ on } (1^{720}0^{720}1^{720}0^{720}0^{720}1^{720}0^{720}0^{720}1^{720})) \\
 & = \langle 0, \frac{7}{8} \rangle \left(\frac{3}{8} \right) \text{ on } \sim \langle -3600, -3600 \rangle (1) \text{ on } \sim \langle -400, 480 \rangle \left(\frac{4}{9} \right) = \langle -2000, -\frac{20153}{18} \rangle \left(\frac{1}{6} \right)
 \end{aligned}$$

- ▶ minimal delay and buffer :

	delay	buffer size
exact result	9 598 (\approx time to receive 5 HD lines)	192 240 (\approx 267 SD lines)
abstract result	11 995 (\approx time to receive 6 HD lines)	193 079 (\approx 268 SD lines)

Conclusion

Ensuring synchronous and other static properties

- ▶ specify/check logical time as special types
- ▶ initially a dependent type system ; now an ML type system with extension by “Laufer & Odersky”
- ▶ this is the way it is done in the Lucid Sychrone compiler the one of SCADE 6
- ▶ some other properties can be expressed as dedicated type-systems (correct initialization of registers, causality analysis)

DSL embedding

- ▶ achieving the same result by designing a DSL (e.g., in Haskell) is difficult
- ▶ how to ensure synchrony, the absence of causality loops, unbounded FIFOs (unless we forbid non-length preserving functions) ?
- ▶ compilation through maximal static expansion does not work well when targeting software code